

TIME DOMAIN STABILITY MARGIN ASSESSMENT OF THE NASA SPACE LAUNCH SYSTEM GN&C DESIGN FOR EXPLORATION MISSION ONE

Keith Clements*

John Wall†

The baseline stability margins for NASA's Space Launch System (SLS) launch vehicle were generated via the classical approach of linearizing the system equations of motion and determining the gain and phase margins from the resulting frequency domain model. To improve the fidelity of the classical methods, the linear frequency domain approach can be extended by replacing static, memoryless nonlinearities with describing functions. This technique, however, does not address the time varying nature of the dynamics of a launch vehicle in flight. An alternative technique for the evaluation of the stability of the nonlinear launch vehicle dynamics along its trajectory is to incrementally adjust the gain and/or time delay in the time domain simulation until the system exhibits unstable behavior. This technique has the added benefit of providing a direct comparison between the time domain and frequency domain tools in support of simulation validation.

INTRODUCTION

The baseline stability margins for SLS were determined by following the well-established classical approach of calculating the gain and phase margins in the frequency domain with a linearized system model. Stability margins derived using the classical methods have been compared against several frequency domain tools, yet only limited efforts have sought to verify these margins in the time domain, the goal of the work as presented herein. This work began as part of the joint SLS and NASA Engineering Safety Center (NESC) effort to demonstrate the stability and flight readiness of the SLS Adaptive Augmenting Control (AAC) system as baselined for the first test flight¹.

Due to the time-varying nature of the gain and phase margins, the nonlinear system was linearized at each time point in order to produce the Nichols charts to determine the stability margins. This was done using the NASA MSFC Frequency Response Analysis and Comparison Tool Assuming Linearity (FRACTAL²) at every second for the duration of SLS's ascent flight. In this paper, the low frequency gain margin will be referred to as the *aero* gain margin (despite its aerodynamic influence being limited to the boost phase of flight), and the high frequency gain margin will be referred to as the *rigid body* gain margin (since it is usually driven by rigid body dynamics).

To determine the amount of gain or phase modification that causes the system to become unstable, the gain or phase of the system was artificially adjusted starting at a pre-determined flight time, and the simulation was run to completion. This was repeated with incrementally increased gain or phase magnitudes until the subsequent response of the axis under investigation showed divergent behavior. Three different time points were investigated, corresponding to each of the flight phases:

* Control System Engineer, NASA MSFC EV41, ERC Inc. - Jacobs ESSSA Contract, keith.r.clements@nasa.gov

† Control System Engineer, NASA MSFC EV41, Dynamic Concepts, Inc. - ESSSA Contract, john.w.wall@nasa.gov

boost, pre-LAS jettison, and post-LAS jettison. The time domain simulation manipulations were performed in the Marshall Aerospace Vehicle Representation in C (MAVERIC) simulation and compared against the classically-derived stability margins using FRACTAL.

A major factor that affects the apparent time domain stability of the vehicle in comparison to the frequency domain margins is the modeling of slosh dynamics. A combination of “requirement” vs. “predicted” damping values and linear vs. nonlinear slosh models were identified as the primary differences and are discussed in this paper.

In the process of determining the aero gain margins, it was seen that the presence of closed-loop guidance in the phase of flight following booster separation (core phase) demonstrated a loss of low frequency gain margin, an effect not predicted by the frequency domain tools. In order to more directly compare the aero margins in these cases to the linear predictions and to ascertain the exact amount of margin degradation due to guidance, the guidance commands were “frozen” (i.e. the last guidance command before modifying the gain was used) for the remainder of flight.

STABILITY MARGIN ASSESSMENT METHOD

The gain or phase adjustment was implemented by applying the axis specific adjustment to the angular acceleration command vector within the optimal control allocator (OCA) as shown in Figure 1. This implementation ensures that the adjustment occurs after the Flight Control System (FCS) internal feedback loop to the translational disturbance compensation algorithm (DCA), thereby ensuring a true total loop modification. This gain or phase application point corresponds to the stability assessment loop break for each axis.

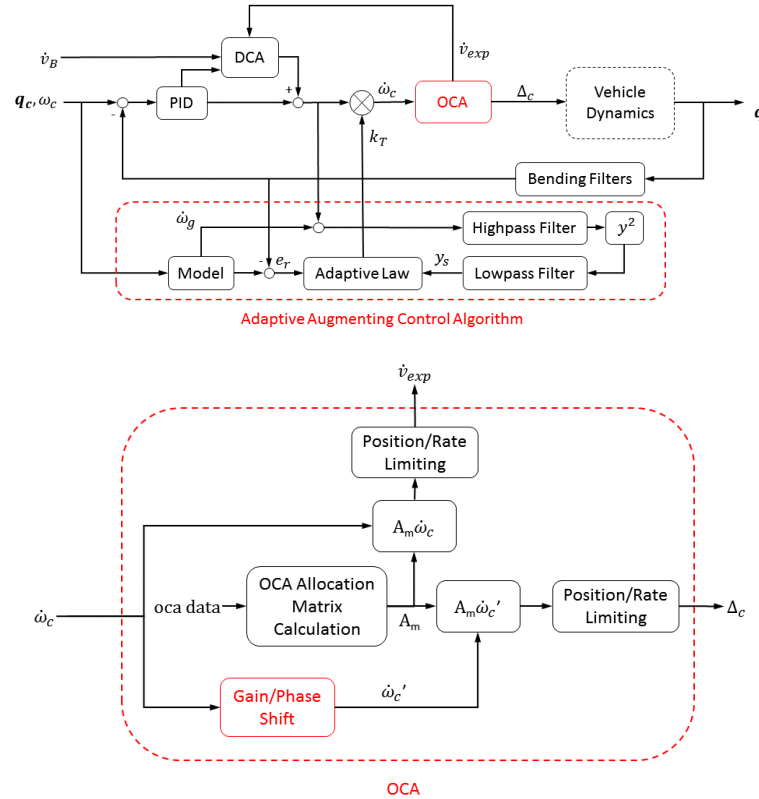


Figure 1: Overall Block Diagram of SLS and OCA Block with Gain/ Phase Adjustment.

Rigid Body Gain Modification Method

Using the time history of stability margins derived from the baseline linearization approach, the time domain derived stability margins were determined via multiple time domain simulations in which axis-specific incremental gain adjustments were made to the nominal system about the expected neutral stability point. The baseline stability margin time histories were used to shift the gain *up* to the frequency-domain values around the zero margin point such that a precise amount of expected instability was maintained throughout flight. When assessing the gain margins, the gain was applied starting at the time point under consideration, thereafter following the variation in the margin found in the linear analysis. Figure 2 shows the baseline margins in FRACTAL with an example of the incremental gain application relative to the FRACTAL margins.

Aero Gain Modification Method

The baseline stability margin time histories were used to shift the gain *down* to various values around the zero margin point such that a precise amount of expected instability was maintained throughout flight. When assessing the aero gain margins, the gain was applied starting at the time point under consideration, thereafter following the variation in the margin found in the linear analysis. Figure 3 shows the baseline stability margins derived from FRACTAL and examples of the incremental gain applications. Due to the quickly varying nature of the boost phase aero gain (up to ~130 seconds) margin and the much slower time constant of the corresponding instability, initial attempts at identifying the aero margin in boost phase proved inconclusive. Only aero gain margins in the core stage of flight after SRB separation are documented herein. Note that in this exo-atmospheric region of flight, the aero margin not actually defined by aerodynamics but by other more constant system dynamics and so is sometimes referred to as the “low frequency gain margin.”




Figure 2: Rigid Body Gain Margins and Modifications over Time.




Figure 3: Aero Gain Margins and Modifications over Time.

Phase Margin Modification Method

When assessing the rigid-body phase margin, a constant time delay was applied to the system starting at the time point under consideration. This delay was converted to a phase in degrees by the formula:

$$\varphi_{deg} = 360 * \omega_{crossover} * t_{delay},$$

where $\omega_{crossover}$ is the 0 dB gain crossover frequency in Hz for the time point under investigation and t_{delay} is the time delay in seconds. Since the linearized phase margins vary in time, the constant delay increment leads to variable instantaneous stability of the system. Overall, it's desirable for the phase margin to be relatively constant to give

the most accurate phase margin at the time point under investigation. Figure 4 shows an example of the phase margins derived in FRACTAL and corresponding phase degradations.

Time Points Analyzed

SLS ascent flight is broken into three main flight phases: boost flight (SRBs and core engines firing), core phase pre-Launch Abort System (LAS) jettison, and core phase post-LAS jettison. The time points for this analysis were chosen to assess the three main phases of flight. These operating points were 80 seconds (boost stage flight), 140 seconds (pre-LAS jettison), and 300 seconds (core stage flight after LAS jettison).




Figure 4: Phase Margins and Modifications over Time.

For each time point, the gain was varied in 0.2 dB steps, and the delay was applied in 0.02 second increments corresponding to a single 50 Hz frame delay. For each time point, the stability margin was found through a numerical assessment of the time domain responses. The metrics and method for assessment of instability depends upon the type of stability margin being verified and the configuration of the system.

Variables Assessed

In order to determine the first simulation that displays unstable behavior, a set of variables was created to fully assess the instabilities. The legend of the gain margin figures shown below show the magnitude of instability, that is, the amount beyond the gain margin for the time point under consideration. This format will be used for the gain margin differences throughout this report.

Body Rates (p, q, or r): These are the main indicators of instability. If the axis under investigation displays divergent oscillations in its body rate, it is said to be unstable. Figure 5 shows an example of the roll rate during an instability applied to the x-axis. It is clear that the first unstable case is 0.2 dB above neutral stability as shown by the diverging roll rate.

Max engine saturation ratio: This variable measures the maximum of the ratio of the commanded gimbal angles to their individual respective limits. If the max engine saturation ratio is larger than one, the system has reached gimbal angle command saturation. For cases with gimbal limits enabled, a value of one is a fair indicator of an ensuing instability. For cases with gimbal limits off, divergent behavior of this variable is indicative of an instability. Figure 6 shows an example of the max engine saturation ratio variable that corresponds to the instability in Figure 5. The frequency content at the onset of the 300 second gain adjustment results from the small magnitude in-flight programmed test input (PTI). Thereafter the slow and small magnitude growth of the max engine saturation ratio corresponds to the case with unstable gain adjustment.

Actuator Duty Cycle: This is the time integral of the absolute value of the pitch/yaw actuator rates. When this variable is observed to quickly diverge, an instability is likely in the system. The added advantage of using the duty cycle is that all values are positive and the magnitudes are that such that inspection of the successive time histories makes the first unstable simulation readily apparent. Figure 7 shows an example of duty cycle responses. The Figure clearly demonstrates that

a gain adjustment of 0.2 dB above neutral stability results in a diverging duty cycle whereas the lower simulations level off fairly soon after the gain modification starts.




Figure 5: Roll Rate for Rigid Body Gain Adjustments Applied at 300 Seconds




Figure 6: Max Engine Saturation Ratio for Rigid Body Gain Adjustments Applied at 300 Seconds




Figure 7: Duty Cycle for Rigid Body Gain Adjustments Applied at 300 Seconds

RESULTS

The method described in the previous section was applied to the specified time points to determine stability margins in the time domain. Figure 8, Figure 10, and Figure 11 show the rigid body, phase, and aero gain margins, respectively, in the time domain. The solid lines represent the margins derived in FRACTAL, and the symbols at specified time points represent the time-domain derived margins.

Figure 8 shows that the rigid body gain margins derived in the time and frequency domain agree within about 0.5 dB for the 80 and 140 second time points. The 300 second time point is unique in the fact that it quickly becomes dominated by slosh dynamics. Regardless of any modeling differences, the instabilities associated with slosh will require a large amount of time to show up in the time domain (~50 seconds). Because of this slow behavior, the margins evaluated in the time domain will often be higher than those derived in the frequency domain. An example of this is shown in Figure 9. The pitch rate of the simulations modified by less than 3.4 dB stay stable while slosh dynamics drive the gain margin; however, a second batch of simulations go unstable at 400 seconds when the rigid body dynamics begin to dominate the margin. This phenomenon associated with slosh will be discussed more in the next section.

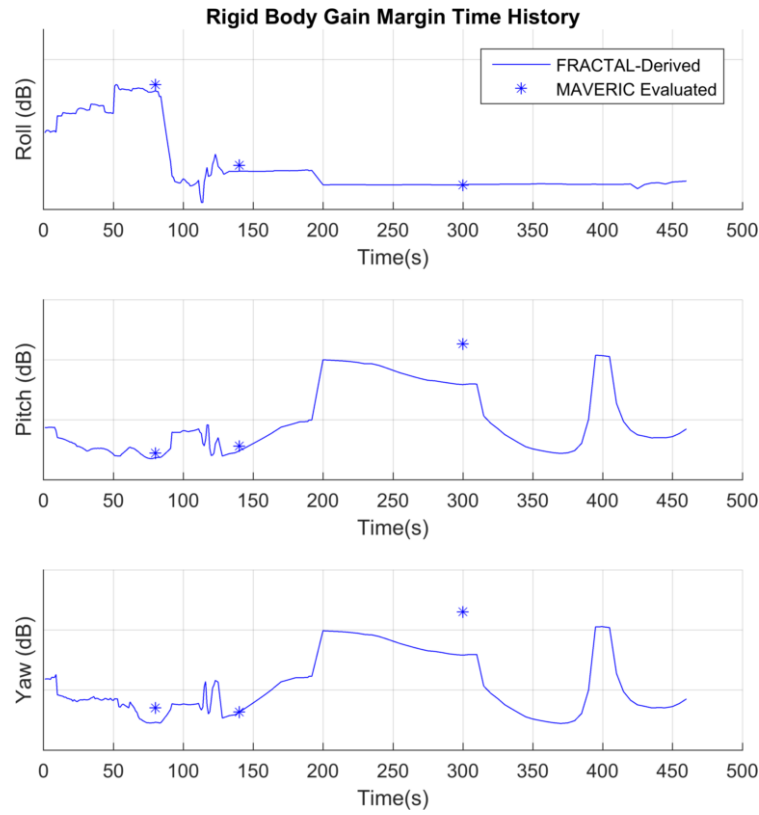


Figure 8: Time-Domain Derived Rigid Body Gain Margins Relative to FRACTAL Margins.




Figure 9: Pitch rate for Rigid Body Gain Modifications Applied at 300 Seconds

Figure 10 shows the rigid phase margins in the MAVERIC time domain. These points have error bars associated with them reflective of the achievable precision in which the phase margin instability can be determined. In addition to the difficulty in instability determination as represented in the error bars, the phase margins are expected to differ for two known reasons. Firstly, during boost phase, the phase margins are changing very quickly. Because it takes at least 10-20 seconds for a phase instability to be realized in the time domain, the margins calculated in the time domain

can be significantly different from those calculated in the frequency domain as time progresses past the point of phase modification. Secondly, the margins calculated in the time domain are overall larger than those calculated in the frequency domain. This is due phase modeling differences between MAVERIC and FRACTAL which yields FRACTAL with a more conservative result. Other sources of discrepancy between the phase margin determination methods are the subject of future study.

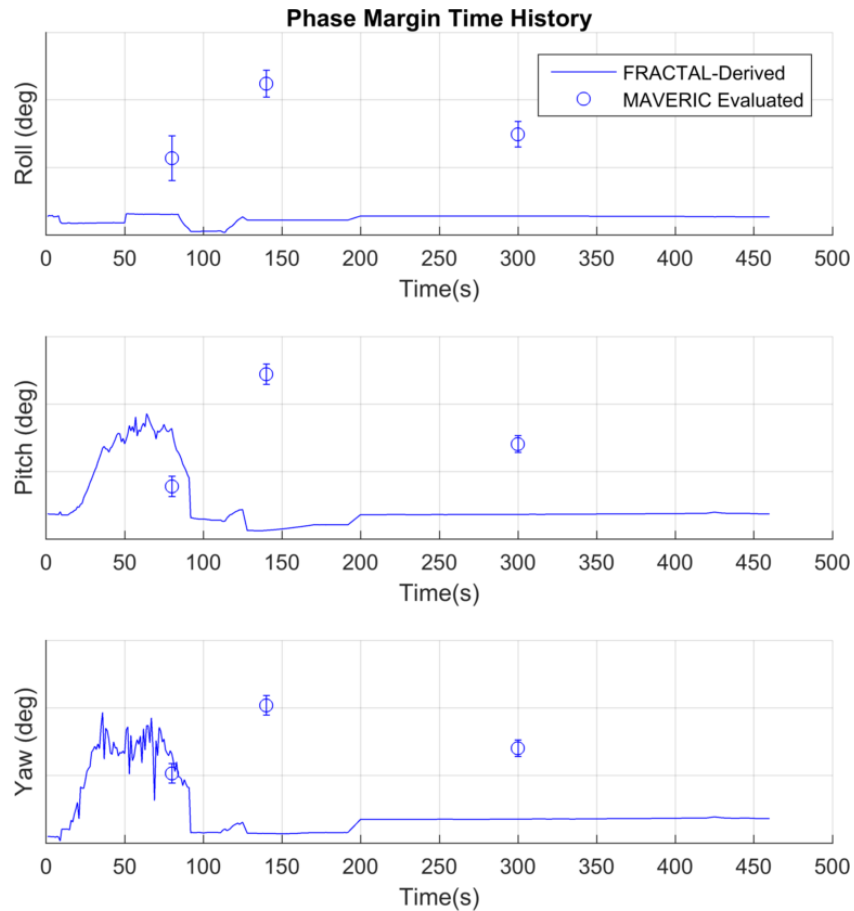


Figure 10: Time-Domain Derived Phase Margins Relative to FRACTAL Margins.

Figure 11 below shows the aero gain margins in the time domain. In this plot, the markers that are below the FRACTAL-derived margins perform better than predicted. In other words, the gain must be lowered beyond the value derived in FRACTAL in order for unstable behavior to be observed.

It was observed that lowering the gain to near the aero gain margin resulted in instabilities earlier (higher) than the margin derived in the frequency domain. It was determined that the closed-loop guidance algorithm was going unstable before the control system. Because of this, the guidance commands were “frozen” (the last command before the gain modifications were applied was used). The results in Figure 11 have frozen guidance commands demonstrating a strong agreement to the FRACTAL predictions. The following section will discuss the difference between regular and frozen guidance.

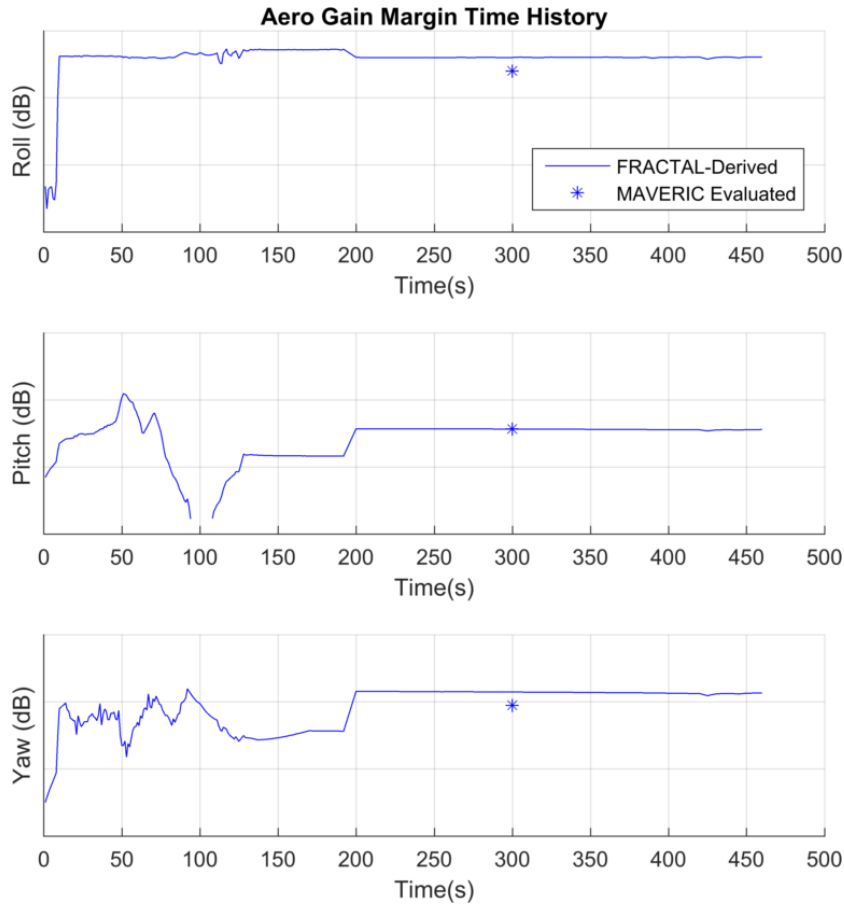


Figure 11: Time-Domain Derived Aero Gain Margins Relative to FRACTAL Margins.

SIGNIFICANT FINDINGS

Frozen Guidance Study

In the process of determining the aero gain margins, it was seen that the presence of closed-loop guidance in core stage flight demonstrated a modest loss of low frequency gain margin, an effect not predicted by the frequency domain tools. In order to more directly compare the aero margins in these cases to the linear predictions, the guidance commands were “frozen” (i.e. the last guidance command before modifying the gain was used) for the remainder of flight. The frozen simulations, as shown in Figure 11, agree exceptional well with the FRACTAL predictions of aero margin.

The method presented in this paper allows for a quantification of the amount of aero gain margin lost due to guidance destabilization. Table 1 shows the aero gain margins with and without closed-loop guidance. It can be seen that the cases with guidance frozen can recover about 3-4 dB more aero gain margin than cases that use closed-loop guidance in pitch/yaw. The roll axis is less sensitive to the presence of the guidance outer loop. **Error! Reference source not found.**

Table 1: Frequency-Relative Aero Gain Margin with and without Guidance Freeze at 300 Seconds.

Axis	Regular Guidance RB Gain Margin Diff From FRACTAL	Frozen Guidance RB Gain Margin Diff From FRACTAL
Roll	-0.5 dB	-0.5 dB
Pitch	4.0 dB	0.0 dB
Yaw	4.0 dB	-0.5 dB

Error! Reference source not found. shows an example of aero gain instabilities at 300 seconds with guidance in its baseline configuration. **Error! Reference source not found.** shows the same example with guidance commands frozen. It can be seen from these two figures with guidance enabled, the onset of system instability occurs at 4 dB higher gain than the same case with frozen guidance commands.




Figure 12: Aero Gain Instability at 300 Second with Regular Guidance.




Figure 13: Aero Gain Instability at 300 Second with Frozen Guidance.

Effect of Slosh Dynamics

With the exception of a brief duration, the slosh dynamics are the driving force behind the gain margin after 300 seconds. In order to quantify the difference in the amount of margin recovered with and without slosh dynamics, time points 300 and 360 seconds were tested with slosh disabled. Figure 14 shows the results with and without slosh at 300 seconds (and slosh off only at 360 seconds). For the 300 second time point, cases with slosh enabled exhibited larger gain margin than the values derived in FRACTAL. This is due to the slow nature of slosh instabilities in the time domain as well as the modeling differences in FRACTAL and MAVERIC.

There are two main slosh damping differences between the FRACTAL configuration used in this study and the MAVERIC time domain simulation that will lead to discrepancies in the rigid gain margin values. First, the predicted baffle damping was used for slosh damping in MAVERIC, whereas the FRACTAL results employed herein use the lower and more constant requirement damping values at 4" wave heights. Second, a recently introduced nonlinear slosh damping model as a function of wave height has been employed in MAVERIC. This more accurate representation of the actual slosh damping in each tank leads to a self-stabilizing system with stable limit cycles rather than the divergent instability predicted by the linear system in FRACTAL. For gain perturbations beyond neutral stability predictions (with slosh as the driving factor behind stability), the slosh wave heights will grow until reaching a damping which yields a neutrally stable system. It can be seen from the results in this section that a significant amount of gain beyond the margins derived in FRACTAL is required to produce an unstable scenario in the time domain. With the inclusion of the nonlinear slosh damping profile, the gain must be adjusted to a point where rigid body dynamics drive the gain margin in order to produce an unstable vehicle.




Figure 14: Time Domain Stability Margins with Slosh Enabled and Disabled

CONCLUSIONS

This paper has defined a method for analyzing the stability margins in the time domain and applied this method to the SLS vehicle to seek corroboration of the stability margins derived in the frequency domain. The margins derived using time and frequency domain techniques agree within expectations, with differences in slosh delay and guidance modeling identified. The slosh damping differences cause the time domain results to recover more rigid body gain margin than predicted in the frequency domain. The rigid body phase margins evaluated in the time domain are consistently higher than those derived in the frequency domain, owing in part to modeling differences and the difficulty in detecting the onset of the slow instability. The aero gain margins were evaluated with and without closed-loop guidance, which revealed that the vehicle displayed unstable behavior about 3-4 dB before the values predicted in the frequency domain.

The method described herein has also been applied the adaptive controller used by SLS. By design, the addition of the adaptive controller can add up to 6 dB of gain margin to the vehicle depending on the nature of the system instabilities, so quantifying the gain benefits in the time domain is very valuable. Another effect that has been explored using this method is the gimbal limiting on the vehicle. SLS has gimbal angle and rate limits associated with the actuators. These limits are not accounted for in the frequency domain, so the method presented in this paper can be used to quantify the effects of gimbal limiting – especially in the presence of an adaptive controller. These two phenomena can be included in any future work on this topic.

REFERENCES

1. VanZwieten, T., et al., “Stability of the SLS FCS with Adaptive Augmentation,” NESC-RP-14-00964, July 2016.
2. Orr, Jeb 2011, “A Flight Dynamics Model for a Multi-Actuated Flexible Rocket Vehicle,” AIAA Atmospheric Flight Mechanics Conference; 8-11 Aug. 2011; Portland, OR; United States

ACRONYMS

AAC	Adaptive Augmenting Control
DCA	Disturbance Compensation Algorithm
FCS	Flight Control System
FRACTAL	Frequency Response Analysis and Comparison Tool Assuming Linearity
LAS	Launch Abort System
MAVERIC	Marshall Aerospace VEHICLE Representation in C
MSFC	Marshall Space Flight Center
NESC	NASA Engineering and Safety Center
OCA	Optimal Control Allocator
PTI	Programmable Test Inputs
RB	Rigid Body
SLS	Space Launch System
SRB	Solid Rocket Booster